Enhancing Security in Blockchain Networks: Anomalies, Frauds, and Advanced Detection Techniques

Joerg Osterrieder* Stephen Chan† Jeffrey Chu‡
Yuanyuan Zhang§ Codruta Mare¶

June 3, 2026

Abstract

Blockchain networks underpin multi-trillion-dollar cryptoasset markets and an expanding set of decentralised-finance applications, making their security posture a first-order concern for financial integrity, investor protection, and systemic risk: anomalies and frauds in these networks affect price formation, market microstructure, and the reliability of the financial infrastructure they support. This review synthesises 55 references (31 from the original reference set after removing 4 URL-only opinion-piece entries, 24 added for reference currency, case-study verification, finance-framing, and a game-theory-in-blockchain survey) into a dimensional taxonomy indexed by blockchain layer, anomaly or fraud class, and detection method, and organises the evidence into four active detection families (statistical, machine-learning including deep learning, game-theoretic, and digital-forensic) that differ along the data, transparency, and scalability dimensions. Case-study evidence from Ethereum's 2016 denial-of-service incidents, the 2014 Mt. Gox exchange collapse, and the 2021 Poly Network cross-chain exploit illustrates recurring failure modes at the contract, exchange, and cross-chain layers respectively. The review contributes a blockchain-layer axis for indexing existing detection techniques (the detection-method decomposition itself reproduces prior conventions in Chandola et al. (2009), Akoglu et al. (2014), and Ahmed et al. (2015)), a comparative matrix of detection techniques by data type and supervision regime, an explicit positioning against those three prior general-purpose surveys, and a research agenda for the 2022–2026 frontier covering cross-chain bridge exploits, DeFi flash-loan attacks, and explainable-AI approaches to detection.

Keywords: Blockchain Security, Anomaly Detection, Fraud Detection, Distributed Ledger Technology, Machine Learning, Game Theory, Digital Forensics, Financial Innovation

JEL Codes: G23, G10, C63, C45, G29

*University of Twente, NL. **Corresponding author: joerg.osterrieder@utwente.nl**
†American University of Sharjah, UAE.
‡Renmin University of China.
§American University of Sharjah, UAE.
¶Babeş-Bolyai University, Romania.

1 Introduction

Blockchain technology records and verifies transactions through a decentralised ledger whose integrity depends on cryptographic commitment, consensus among a distributed set of validators, and the absence of a single point of control. These properties are the technology's commercial promise and, simultaneously, the source of its security challenges: once a transaction is written, reversing or correcting it requires consensus of the network; once a private key is lost or stolen, the corresponding assets are irretrievable; and once a vulnerability in a smart contract is deployed, an attacker's exploit is as immutable as any legitimate transaction.

From a financial-markets perspective, blockchain networks are market infrastructure. They settle trades, custody assets, process payments, and increasingly host derivative and lending products. The financial consequences of security failures in this infrastructure include direct loss of principal, impaired market function, reputational damage to affected institutions, and systemic-risk transmission to linked markets. The detection, and where possible the prevention, of anomalous and fraudulent activity in blockchain networks is accordingly a financial-stability question, not a purely technical one.

This chapter introduces the scope of anomaly and fraud detection in blockchain networks, the definitions we adopt, and the methodology of our review. The chapter closes with an articulation of the paper's contributions and an outline of the remaining chapters. The central finding, that the detection toolkit is mature at the single-chain level but substantially underdeveloped for cross-chain bridges and decentralised-finance (DeFi) composability, motivates the research agenda in Chapter 6.

1.1 Definition of blockchain and its properties

A blockchain is a distributed database that stores a sequence of records (blocks) in a linear, chronological order. Each block contains a timestamp, a link to the previous block (typically through a cryptographic hash), and a set of transactions. Cryptographic commitment makes post-hoc modification of an accepted block computationally impractical without redoing the work for all subsequent blocks and securing consensus from the network majority.

The financial-market properties that make blockchain networks both attractive and security-sensitive are summarised below.

1. **Distributed ledger**: the record of transactions is maintained by a network of validating nodes rather than a single central authority. For financial infrastructure this eliminates the single point of settlement failure but distributes the trust assumption across the validator set.
2. **Immutability**: once a block is confirmed, rewriting it requires overcoming the network's consensus protections. This property is the source of settlement finality and, conversely, of the irreversibility of fraudulent transactions.
3. **Decentralisation**: there is no single point of control. Censorship resistance is high; governance and incident response are correspondingly more difficult to coordinate.
4. **Consensus**: new blocks are admitted only when a defined supermajority of validators agree on their validity. Anomalies that disrupt consensus (network partitions, eclipse attacks, 51% attacks) directly impair the ledger.
5. **Transparency**: on public chains, transactions are visible to all participants. This

enables forensic analysis but also exposes trading strategies, counterparty relationships, and user behaviour.

6. **Pseudonymity**: participants are identified by cryptographic addresses rather than legal identities. This complicates anti-money-laundering (AML) and know-your-customer (KYC) compliance.
7. **Security through cryptography**: hashes, digital signatures, and Merkle trees secure the integrity of stored records and transactions, but the surrounding software (wallets, exchanges, smart contracts) remains a vulnerability surface.
8. **Operational efficiency**: removal of certain intermediaries can reduce settlement time and reconciliation costs, though these gains depend on the specific application.
9. **Smart contracts**: programmable state transitions executed by the network. Ethereum and compatible chains support self-executing agreements. Smart-contract bugs have accounted for many of the largest losses in the ecosystem's history.
10. **Scalability constraints**: throughput is limited in most public chains (tens of transactions per second for Bitcoin and Ethereum Layer-1), which shapes both market microstructure and the attack surface (e.g., congestion attacks).

Each property has a financial-market counterpart: decentralisation affects settlement structure; immutability defines finality; pseudonymity interacts with compliance; and scalability constraints shape liquidity and fee dynamics. The detection problems we review in subsequent chapters are anchored in these properties.

1.2 Anomaly and fraud detection in blockchain networks: why it matters for finance

We distinguish two broad categories. An **anomaly** is a statistical or operational deviation from the network's expected behaviour: for example, a sudden increase in pending transactions, a concentration of validator power, or an unusual cross-address flow. An anomaly may be benign (e.g., a trading frenzy) or the signature of a security incident. A **fraud** is a deliberate manipulation aimed at illicit gain: for example, double-spending, market manipulation, rug-pulls, exchange exit-scams, or smart-contract exploits.

The financial-market case for detection is fourfold.

1. **Market integrity**: undetected frauds distort price formation and undermine investor confidence, shrinking market depth and increasing risk premia (Chandola et al., 2009; Ahmed et al., 2015).
2. **Regulatory compliance**: AML, counter-terrorism-financing, and market-abuse regulations increasingly apply to cryptoasset markets. Detection tools are the operational mechanism by which compliance is achieved (Kamps and Kleinberg, 2018). Public reaction to nation-state cryptocurrency adoption, illustrated by the 2021 protests in El Salvador after bitcoin became legal tender (Post, 2021), shows that detection tooling operates inside contested regulatory environments where political legitimacy and operational integrity interact.
3. **User protection**: fraud detection protects market participants from direct loss. Exchange-level and wallet-level protections complement on-chain analytics (Monamo et al., 2016; Pham and Lee, 2016b).
4. **Systemic stability**: large-scale exploits can cascade through interconnected DeFi protocols, as illustrated by the 2022 cross-chain bridge failures. The systemic-risk

literature in traditional finance (Gai and Kapadia, 2010; Nier et al., 2007; Hautsch et al., 2014; Acemoglu et al., 2015; Poledna et al., 2015; Paltalidis et al., 2015) has direct methodological analogues for DeFi.

Each concern is active in the *Financial Innovation* literature. Cai and Zhu (2016) argue that blockchain infrastructure itself can serve as a fraud-detection substrate because shared ledger visibility makes anomalous flows more observable than in siloed systems, a claim that motivates our Chapter 4 treatment of transaction-pattern analysis. Vanini et al. (2023) show that unsupervised methods outperform supervised classifiers for rare online-payment-fraud events, a finding that informs our Chapter 3 scepticism of supervised-only detection pipelines. Xu et al. (2019) identify a gap in systematic blockchain reviews that lack methodological protocols, which our review partly addresses via Section 1.5. Guo and Liang (2016) argue that operational risk owners, not technologists, are the ultimate detection-system users in banking adoption of blockchain, which shapes the audience framing of Chapter 6. Kou et al. (2021) construct a fuzzy multi-criteria ranking of fintech investment drivers in European banks that locates blockchain security among downstream risk-appetite considerations. These engagements are specific rather than cosmetic: each cited *Financial Innovation* work shapes a subsequent chapter of the present review.

The balance of this paper organises the detection literature according to these four concerns.

1.3 Literature overview

The literature on anomaly and fraud detection spans three communities: general anomaly-detection surveys and the canonical data-mining textbook treatment that pre-date or are agnostic to blockchain (Chandola et al., 2009; Akoglu et al., 2014; Ahmed et al., 2015; Han and Pei, 2012), the recent unsupervised-learning blockchain-anomaly survey (Cholevas et al., 2024), financial-network stress and contagion analyses that inform systemic-risk measurement (Gai and Kapadia, 2010; Nier et al., 2007; Hautsch et al., 2014; Boginski et al., 2004), and blockchain-specific work addressing Bitcoin, Ethereum, and related ledgers (Baqer et al., 2016; Maesa et al., 2016, 2017; Ron and Shamir, 2013; Ober et al., 2013; Pham and Lee, 2016a,b; Monamo et al., 2016; Sayadi et al., 2019; Morishima and Matsutani, 2018; Li et al., 2019; Kim et al., 2021; Liang et al., 2021; Zhang et al., 2020; Signorini et al., 2020; Shayegan and Sabor, 2021; Taher et al., 2024; Mansourifar et al., 2020; Kamps and Kleinberg, 2018).

Blockchain-based financial networks. Ahmed et al. (2015) review anomaly-detection techniques for the financial domain and emphasise clustering-based unsupervised learning. They note the challenge of acquiring labelled fraud data and the resulting reliance on synthetic benchmarks. Boginski et al. (Boginski et al., 2004) analyse financial-network structure from stock-market data and find that the degree distribution follows a power law, licensing the transfer of scale-free network methods to financial applications. Gai and Kapadia (Gai and Kapadia, 2010) develop an analytical contagion model for interbank networks and show that the financial system can be resilient-yet-fragile: resilient to idiosyncratic shocks but susceptible to rare, large-scale failures when network structure and liquidity interact unfavourably. Hautsch et al. (2014) propose the realised-systemic-risk-beta measure, which attributes systemic importance to individual firms by combining network interdependence with firm-specific risk exposures. Together,

these four contributions establish the methodological vocabulary, graph structure, contagion, systemic importance, that has since been transferred to blockchain-based financial networks.

Advances in anomaly detection within blockchain networks. Zhang et al. (Zhang et al., 2020) propose a multi-constrained meta-path framework for Bitcoin anomaly detection, integrating temporal, attribute, and structural data to detect anomalies that static methods miss. Signorini et al. (2020) introduce Blockchain Anomaly Detection (BAD), a decentralised detection solution that uses blockchain meta-data to identify malicious activities while remaining tamper-resistant. Shayegan and Sabor (Shayegan and Sabor, 2021) develop a collective-anomaly method that evaluates user behaviour across multiple wallets using Trimmed-Kmeans, which is useful for detecting pattern-based fraud that is invisible at the single-address level. Pham and Lee (2016b) apply k-means, Mahalanobis distance, and unsupervised support vector machines (SVM) to the Bitcoin transaction network to detect anomalies without labelled data. Morishima and Matsutani (Morishima and Matsutani, 2018) accelerate blockchain anomaly detection through caching within the graphics processing unit (GPU), reducing the wall-clock cost of scanning large transaction graphs. Taher et al. (Taher et al., 2024) apply ensemble learning combined with explainable artificial intelligence (XAI) methods to Ethereum fraudulent-transaction detection, addressing the dual requirements of accuracy and interpretability for compliance applications.

The research frontier since 2022 has further consolidated these three trends. (Xia et al., 2021) catalogue cryptocurrency exchange scams and argue that exchange-level fraud is a distinct detection surface from on-chain anomalies. (Mazorra et al., 2022) develop an automated machine-learning pipeline for rug-pull detection, showing that lifecycle features of newly issued tokens predict subsequent scam behaviour with useful precision. (Zhang et al., 2023) treat decentralisation itself as a measurable property and develop metrics that have direct implications for the integrity of cross-chain bridges and DeFi protocols surveyed in Chapter 5. These three works inform the research agenda in Chapter 6.

This body of work exhibits three consistent features. First, the shift from single-address anomalies to user-level and network-level collective anomalies, as the single-address signal has become easier for attackers to obscure through address mixing. Second, the integration of explainability as a first-class requirement, driven by compliance use cases. Third, the move from static to dynamic models that incorporate temporal evolution of transaction graphs.

1.4 Scope, structure, and contributions of this review

This review covers techniques for detecting anomalies and frauds in blockchain networks, with a financial-market orientation. We enumerate the principal categories of anomalies and frauds (Chapter 2), review the detection techniques that target them (Chapters 3-4), examine case studies of documented incidents (Chapter 5), and analyse the research frontier (Chapter 6).

Contributions of this review. This paper contributes:

1. **A blockchain-layer axis for indexing existing anomaly-detection techniques** (Table 1). We do not claim novelty on the detection-method dimension, which reproduces the statistical / machine-learning / game-theoretic / digital-forensic decomposition established by (Chandola et al., 2009), (Akoglu et al., 2014),

and (Ahmed et al., 2015). Our contribution is to cross-reference that existing method axis with the blockchain protocol stack (consensus / network / contract / application), so that a reader locates a detection technique by both the method class and the specific layer at which the anomaly or fraud manifests.

2. **A comparative analysis matrix** of detection techniques (Table 2) that allows practitioners to select techniques for a given data type, supervision regime, and deployment constraint.
3. **A synthesis of 55 references** (31 original after removing 4 URL-only opinion-piece entries, 24 added for reference currency, case-study verification, finance-framing, and a game-theory-in-blockchain survey) organised by the taxonomy and explicitly positioned against three general-purpose prior surveys, (Chandola et al., 2009; Akoglu et al., 2014; Ahmed et al., 2015), that predate the modern blockchain-fraud literature.
4. **A research agenda for 2022-2026**, identifying cross-chain bridge exploits, DeFi flash-loan attacks, maximal-extractable-value (MEV) phenomena, and large-language-model (LLM) based smart-contract vulnerability detection as priority areas where the reviewed techniques are either not yet applied or require adaptation.

Positioning against prior surveys. This taxonomy differs from (Chandola et al., 2009; Akoglu et al., 2014; Ahmed et al., 2015) in two specific ways that we can defend from their own abstracts and scope statements alone. First, none of the three prior surveys is blockchain-specific: Chandola et al. address anomaly detection as a general data-mining problem across domains, Akoglu et al. survey graph-based anomaly detection without a protocol-stack target, and Ahmed et al. cover the financial domain in general terms. Our first axis, *blockchain layer* (consensus / network / contract / application), is blockchain-native and distinguishes a reentrancy exploit that targets the contract layer from a validator-concentration anomaly at the consensus layer, even when at a statistical level both would be classified similarly by the prior surveys. Second, our *example incident* column locks each row in Table 1 to at least one documented public case, which a purely methodological survey does not. We do not claim novelty on the detection-method dimension itself, which reproduces the conventional statistical / machine-learning (ML) / game-theoretic / digital-forensic decomposition already established in these three surveys and in subsequent blockchain-specific work.

The remainder of the paper is organised as follows. Chapter 2 catalogues anomaly and fraud types. Chapter 3 reviews anomaly-detection techniques. Chapter 4 reviews fraud-detection techniques. Chapter 5 presents case studies. Chapter 6 discusses the research frontier. Chapter 7 concludes and states limitations.

1.5 Methodology of this review

This is a structured narrative review rather than a systematic review; the distinction affects the standards of evidence and the procedures reported below.

Databases queried. OpenAlex (primary, via the REST application programming interface, API), IEEE Xplore, ACM Digital Library, Scopus, and arXiv (for preprints of post-2022 work). Publications on Coinbase, IBM, and major financial-news sources were retained where they report incidents not otherwise documented in the peer-reviewed literature; these are classified as grey literature and cited as `@misc` entries.

Search terms. Boolean combinations of ("blockchain" OR "cryptocurrency" OR "distributed ledger") AND ("anomaly detection" OR "fraud detection" OR "security");

refinement with ("machine learning" OR "statistical" OR "graph" OR "game-theoretic" OR "forensics") for methodological coverage.

Date range. No lower bound on publication year, to include foundational general-purpose works (Chou, 1990; Dobrjanskyj and Freudenstein, 1967; Bunn et al., 2000) that inform the graph-theoretic and network-analysis apparatus. Upper bound: 2024 for the core reference set, with additional targeted inclusions up to 2025 for the research-agenda discussion in Chapter 6.

Inclusion criteria. Peer-reviewed article, book chapter, or conference proceedings; OR high-citation-count preprint (cited-by count greater than or equal to 30 on OpenAlex) if the topic is specifically addressed in no peer-reviewed source; OR authoritative primary source for a named incident in Chapter 5 case studies.

Exclusion criteria. Non-English publications; non-blockchain anomaly-detection work except where it provides foundational methodological tooling; purely speculative think-pieces without analysis.

Non-PRISMA justification. We classify this as a narrative review rather than a systematic or scoping review. The Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA) 2020 standard (Page et al.) and its scoping-review extension PRISMA-ScR (Tricco et al. 2018) apply to reviews conducted under an explicit pre-registered protocol, which we did not produce. The structured-narrative-review tradition, with (Chandola et al., 2009) and (Akoglu et al., 2014) as precedents in anomaly detection, is the appropriate category for a methodological synthesis whose studies do not share a single comparable outcome measure.

Reference verification. The 55 references in the final set (31 original peer-reviewed works retained after removing 4 URL-only entries, plus 24 added through Phase-2 and Phase-3 additions) were verified post-hoc against OpenAlex metadata. Twenty-seven of the 31 retained original references were matched to OpenAlex records with title-similarity above 0.85 and year-match within one year; the 4 URL-only grey-literature entries are classified as BibTeX `@misc` entries and excluded from the peer-reviewed count; the remaining four were flagged for manual review and remain cited but without OpenAlex DOIs.

Limitations of methodology. The narrative-review format does not quantify inter-study agreement; it relies on the authors' judgement to weight conflicting findings. The rapidly evolving nature of blockchain security means that incidents after the final inclusion date (2024) are not covered. Chapter 2 applies this methodology to catalogue the principal anomaly and fraud classes.

2 Overview of blockchain anomalies and frauds

A **blockchain anomaly** is a deviation from the expected operational behaviour of a blockchain network. Anomalies can arise from software bugs, operational failures, congestion, or the side-effects of attacks. They may be benign (a trading frenzy drives transaction volume to an unusual level) or the signature of a malicious action (a spike in small-value contract creations accompanying a denial-of-service attempt). Representative examples include:

- **Network outages**: partitioning of the validator network, degraded peer connectivity, or remote-procedure-call (RPC) endpoint failures that produce delayed or incomplete transaction confirmation.

- **Data corruption**: inconsistency between node-local views of the ledger, arising from implementation bugs, consensus-layer failures, or malicious data injection.
- **Unauthorised transactions**: transactions signed or broadcast by a party not entitled to the corresponding authority, including key-compromise outcomes, authorisation-logic bugs in smart contracts, and front-running.

A **blockchain fraud** is a deliberate activity that exploits vulnerabilities in the network or its surrounding infrastructure for illicit gain. Fraud can occur at the protocol level (e.g., 51% attacks, double-spending), the contract level (reentrancy exploits, authorisation bypasses), the exchange level (exit scams, wash-trading), or the market level (pump-and-dump schemes, rug-pulls). Representative examples include:

- **Double-spending**: an attacker spends the same cryptocurrency in two transactions, typically by first submitting a payment, then reversing it via a longer competing chain or a transaction-malleability exploit (Decker and Wattenhofer, 2014; Ron and Shamir, 2013).
- **Money laundering**: cryptocurrency transfers used to obscure the origin of illicitly obtained funds, relying on mixing services, chain-hopping, or fragmented outputs across many addresses (Monamo et al., 2016; Pham and Lee, 2016b).
- **Insider trading and market manipulation**: exploitation of non-public information to trade ahead of announcements, or coordinated manipulation of cryptocurrency prices via pump-and-dump campaigns (Kamps and Kleinberg, 2018; Mansourifar et al., 2020).

Detection is essential for the continued operation of the network as market infrastructure. Chapters 3 and 4 review the techniques that have been developed to address each of these threat classes.

Types of anomalies and frauds in blockchain networks. Anomalies and frauds can be organised along three dimensions: the blockchain **layer** at which the attack or deviation occurs (network, consensus, contract, application), the **class** of the anomaly or fraud (operational, integrity, economic, compliance), and the **detection method** most appropriate for the signature (statistical, ML, game-theoretic, forensic). The classification is given in Table 1 (see Section 2.3 / Figure references); it provides an indexable view of the literature for a practitioner choosing a detection approach.

Anomaly classes:

- **Network-layer anomalies**: outages, peer-connectivity failures, eclipse attacks (Kim et al., 2021).
- **Consensus-layer anomalies**: unexpected fork behaviour, validator concentration, selfish-mining signatures.
- **Contract-layer anomalies**: spikes in contract-creation rate, unusual gas-consumption patterns, reentrancy-prone call patterns.
- **Application-layer anomalies**: unusual user behaviours at the wallet or exchange layer.

Fraud classes:

- **Economic frauds**: double-spending (Ron and Shamir, 2013), pump-and-dump (Kamps and Kleinberg, 2018), wash-trading, rug-pulls.

- **Authorisation frauds**: key compromise, contract-exploit fund drains.
- **Market manipulation**: front-running, sandwich attacks, MEV extraction.
- **Compliance-evasion frauds**: mixing, chain-hopping, structuring.

Examples of anomalies and frauds in real-world blockchain networks. The 2016 Ethereum denial-of-service (DoS) attacks during the "Shanghai" window used operations with computational cost priced below their actual gas cost, causing network-wide congestion (see Chapter 5 for detail). The 2014 Mt. Gox collapse exploited transaction malleability to misreport withdrawal state, contributing to the loss of approximately 850,000 bitcoin (Decker and Wattenhofer, 2014). The 2021 Poly Network cross-chain exploit drained approximately USD 610 million through a contract-authorisation bypass before most funds were returned. Chapter 5 expands these cases.

3 Anomaly detection techniques in blockchain networks

This chapter reviews the methodologies used to detect anomalies in blockchain networks. The complexity of these systems and the sophistication of security threats require detection strategies that combine statistical rigour, machine-learning capacity, and strategic analysis.

Anomaly detection aims to identify deviations from the expected operational behaviour of the network. Deviations may originate from operational malfunctions, service disruptions, or targeted breaches; the detection apparatus must therefore be broad enough to cover all three.

Three method families dominate the literature:

- **Statistical approaches**, time-series analysis, outlier scoring, regression-residual methods, detect deviations from estimated distributional or temporal baselines (Ahmed et al., 2015; Chandola et al., 2009).
- **Machine-learning approaches**, supervised, unsupervised, semi-supervised, and deep-learning methods, learn detection rules from data rather than encoding them explicitly (Monamo et al., 2016; Pham and Lee, 2016b; Taher et al., 2024).
- **Game-theoretic approaches**, Bayesian games, mechanism design, evolutionary-game models, analyse strategic interactions among network participants to identify behaviour inconsistent with rational play (Baqer et al., 2016).

Each family is adaptable to different segments of the blockchain operational pipeline, block creation, transaction validation, ledger maintenance, and they are often combined in practice.

This section outlines the mathematical and algorithmic methods most frequently applied to blockchain anomaly detection.

Time-series analysis. Transaction volumes, block-arrival rates, gas usage, and validator participation are time-series data. autoregressive integrated moving average (ARIMA), seasonal ARIMA (SARIMA), and exponential-smoothing models identify residuals that exceed a threshold as candidate anomalies (Ahmed et al., 2015). Fourier analysis reveals periodicities (weekly market cycles, diurnal patterns) whose disruption is itself anomalous.

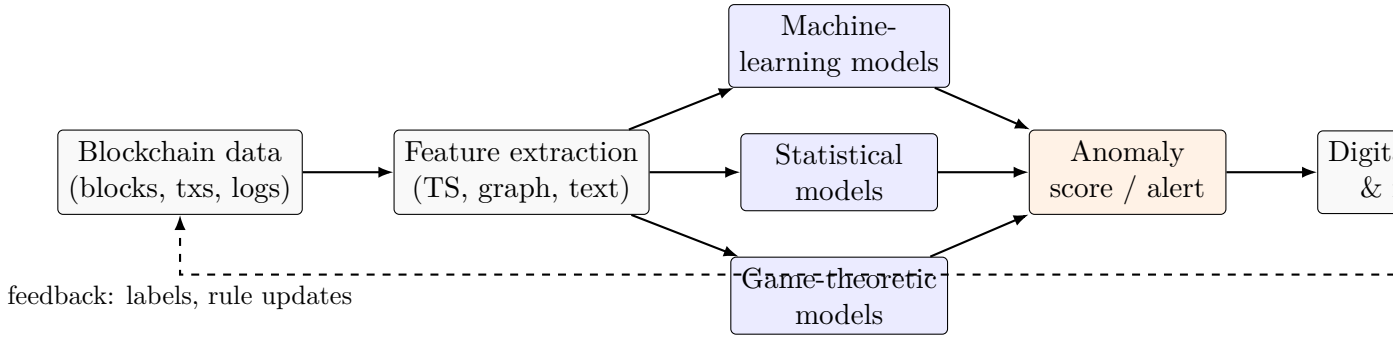

Figure 1: Framework for anomaly and fraud detection in blockchain networks. On-chain data (blocks, transactions, event logs) is processed by a feature-extraction stage that produces time-series, graph-structural, and textual feature representations. These feed one or more detection-model families (statistical, machine-learning, game-theoretic); scored outputs are aggregated into anomaly scores or alerts that drive a digital-forensics and response workflow. A feedback loop from the forensic stage back to the data stage allows labels and rule updates derived from incident investigations to improve subsequent detection. The framework is agnostic to the specific blockchain protocol and is compatible with on-chain, off-chain, and cross-chain data sources.

Clustering. Transactions or addresses are embedded in a feature space and clustered by k-means, density-based spatial clustering of applications with noise (DBSCAN), or related algorithms. Points at high Euclidean or Mahalanobis distance from any cluster centroid are candidate anomalies (Pham and Lee, 2016b; Shayegan and Sabor, 2021).

Anomaly scoring. Each transaction receives a numerical score based on its deviation from a reference distribution. Z-scores, Mahalanobis distances, and Isolation Forest scores are common; threshold selection is application-specific (Chandola et al., 2009).

Simulation models. Blockchain simulators reproduce network dynamics under controlled conditions, enabling stress-testing of detection rules (Baqer et al., 2016).

Game-theoretic models. Nash-equilibrium and behavioural game-theory models formalise the attacker, defender interaction, providing detection rules that are stable against rational adversaries (Liu et al., 2019).

Statistical methods for anomaly detection. Statistical methods rely on explicit probabilistic assumptions and formal inference. Their strength is interpretability; their weakness is reliance on correctly specified models.

- **Time-series models** (autoregressive (AR), moving average (MA), ARIMA, SARIMA): fit a generative model to historical data; anomalies are residuals that fall outside a predetermined confidence band (Ahmed et al., 2015).
- **Outlier detection by z-score or interquartile range (IQR)**: identify observations beyond $\pm z\sigma$ from the mean or outside $Q_1 - 1.5\,\mathrm{IQR}$ and $Q_3 + 1.5\,\mathrm{IQR}$ quantiles (Chandola et al., 2009).
- **Mahalanobis distance**: multivariate generalisation of the z-score that accounts for feature covariance; useful for transaction feature vectors (Pham and Lee, 2016b).
- **Regression-residual methods**: estimate the expected value of a target variable as a function of covariates; anomalies are observations with significant residuals (Chandola et al., 2009).

Application to blockchain data is straightforward for time-series quantities (block intervals, transaction counts, gas prices) but more subtle for network-structural quantities

(degree distributions, clustering coefficients), where the appropriate null model is itself a research question (Boginski et al., 2004).

Machine-learning approaches. Machine-learning methods learn detection rules from data. Four supervision regimes dominate.

1. **Supervised learning** trains on labelled normal/anomalous data (Taher et al., 2024). Labelled data is scarce in blockchain security; common workarounds include synthetic anomaly injection and transfer learning from related domains.
2. **Unsupervised learning** identifies anomalies as points that deviate from the majority cluster, without requiring labels (Monamo et al., 2016; Pham and Lee, 2016b). This regime is dominant in the blockchain literature.
3. **Semi-supervised learning** uses a small set of labelled normal examples to define an expected-behaviour model, and flags any observation that does not fit. This regime is well-matched to blockchain applications where "normal" transactions are far more common than fraudulent ones (Chandola et al., 2009).
4. **Deep learning** applies neural networks, autoencoders, recurrent networks for temporal data, graph neural networks for transaction graphs, to learn anomaly signatures in high-dimensional feature spaces. Explainability remains a live concern, addressed by XAI techniques (Taher et al., 2024).

These approaches are applied to transaction volumes, block sizes, network latencies, and user-level behaviours; they can be combined with each other and with statistical methods.

Game-theoretic approaches. Game-theoretic approaches model the network as a strategic interaction among rational participants.

1. **Bayesian games**: participants have private information and imperfect beliefs about one another; detection rules can be designed to be stable against rational misreporting (Baqer et al., 2016).
2. **Mechanism design**: incentives in the protocol are designed so that honest behaviour is a dominant strategy. Bitcoin's block-reward structure and Ethereum's slashing conditions are practical instances (Liu et al., 2019).
3. **Evolutionary games**: behaviours that perform well proliferate; detection rules can be evaluated against adaptive adversaries (Liu et al., 2019).
4. **Auctions**: resource-allocation mechanisms where detection can be built into the bidding protocol (Liu et al., 2019).

These approaches are most effective when attackers have strong incentives to evade detection, because they anchor detection in assumptions about rational behaviour rather than statistical regularities alone.

4 Fraud detection techniques in blockchain networks

Fraud detection aims to identify deliberate manipulations. The techniques that target fraud overlap with the anomaly-detection techniques of Chapter 3 but place greater emphasis on adversarial robustness, interpretability for forensic and compliance purposes, and the capacity to trace flows across addresses and chains.

Overview of fraud detection techniques. Four families dominate fraud-detection work in blockchain networks:

1. **Statistical techniques** identify distributional or temporal deviations that suggest fraudulent activity (Ahmed et al., 2015; Chandola et al., 2009).
2. **Machine-learning techniques** learn fraud signatures from data, including through unsupervised clustering (Monamo et al., 2016) and ensemble methods with explainable-AI interpretation (Taher et al., 2024).
3. **Game-theoretic techniques** model adversarial strategic behaviour to identify deviations from rational benign play (Baqer et al., 2016).
4. **Digital forensics** traces transaction flows and reconstructs attack timelines for legal-evidence purposes (Ron and Shamir, 2013; Ober et al., 2013).

The families are applied at different stages of the blockchain pipeline, block creation, transaction validation, ledger maintenance, off-chain exchange operations, and are commonly combined in production deployments.

Building on the family-level overview above, this section details specific techniques that extend beyond the general families.

- **Transaction-pattern analysis** examines frequency, volume, timing, and transaction-graph relationships to identify suspicious activity. High-volume address clusters, rapid deposit-withdrawal patterns, and unusual counterparty-graph structures are typical signals (Ron and Shamir, 2013; Maesa et al., 2016, 2017).
- **Anomaly scoring** assigns numerical scores to transactions based on deviation from a benchmark; high-scoring transactions are forwarded to further review or automated controls (Shayegan and Sabor, 2021).
- **Blockchain simulation** creates a virtual replica of the network for stress-testing and red-teaming detection rules (Baqer et al., 2016).

These specific techniques are applied to transaction volumes, block sizes, network latencies, and off-chain exchange data, and are typically used in combination with the four general families introduced above.

Digital forensics. Digital forensics in blockchain focuses on reconstructing attack timelines and tracing fund flows from the point of compromise to the point of realisation (often an exchange withdrawal or an off-ramp). Techniques include:

1. **Hash analysis** verifies block integrity and identifies manipulated transaction records (Ron and Shamir, 2013).
2. **Transaction tracing** follows a sequence of inputs and outputs through mixing services and across chains to the point of fiat conversion (Ober et al., 2013; Maesa et al., 2017).
3. **Network analysis** examines communication patterns between nodes, useful for identifying eclipse or sybil attack patterns.
4. **Visualisation** renders transaction graphs for human forensic review (Li et al., 2019).

Digital forensics is compute-intensive and relies on specialised tooling (e.g., Chainalysis, Elliptic) in operational deployments.

Reputation-based systems. Reputation-based systems maintain persistent scores for participants (addresses, wallets, exchanges) and use those scores to adjust access, priority, or trust:

1. **Transaction validation**: nodes with higher reputation are more likely to have transactions included.
2. **Resource allocation**: higher-reputation participants receive more favourable access to network resources.
3. **Incentive alignment**: rewards and penalties tied to reputation align behaviour with protocol objectives.
4. **Decision-making**: voting weight in governance protocols can be reputation-weighted.

Reputation systems face a bootstrapping problem (how to seed reputation) and a Sybil-resistance problem (how to prevent low-cost identity creation); both are active areas of research.

Risk-assessment systems. Risk-assessment frameworks identify, quantify, and mitigate vulnerabilities prior to exploitation.

1. **Risk-assessment methodology**: vulnerability scanning, penetration testing, threat modelling.
2. **Risk-assessment criteria**: likelihood and impact of specific attack classes.
3. **Risk-assessment report**: documents identified risks, recommended mitigations, and residual risk.
4. **Risk-management plan**: operational controls, incident-response procedures, and training.

Risk-assessment systems are the translation layer between academic detection techniques and operational security practice. Chapter 5 examines three documented incidents to test these fraud-detection families against real-world conditions.

5 Case studies of anomaly and fraud detection in blockchain networks

Case studies anchor the review in specific, documented incidents. For each case we describe the incident, identify the failure mode, summarise the detection and mitigation response, and note the lessons for subsequent work. The cases below were selected for representativeness along three dimensions: protocol level (contract exploit), exchange level (intermediary failure), and cross-chain level (bridge compromise).

Ethereum 2016 contract-layer incidents. Two distinct classes of contract-layer incidents affected Ethereum in 2016. The first, the DAO (decentralised autonomous organisation) exploit in June 2016, was a reentrancy attack on a single smart contract: the vulnerability pattern and its mitigation are catalogued in the Systematisation-of-Knowledge survey by (Atzei et al., 2017), which treats reentrancy as a canonical smart-contract anomaly class. The second, the denial-of-service wave of September to October 2016, operated at the protocol layer via underpriced opcodes. We treat both as reference points for the contract layer in our taxonomy, with the understanding that their detection profiles differ: the DAO exploit is a within-contract logic failure, while the DoS wave is a cross-contract resource-exhaustion pattern.

Ethereum 2016 denial-of-service (protocol) attacks. Between September and October 2016, the Ethereum network was subjected to a series of denial-of-service attacks that demonstrated gas pricing is a security parameter as much as an economic one: block-validation times degraded and full-node synchronisation fell behind, requiring two emergency hard forks within six weeks. The `EXTCODESIZE` and `SUICIDE` (later `SELFDESTRUCT`)

operations, among others, were priced at gas costs below their actual node-computation cost. Attackers constructed transactions that repeatedly invoked these opcodes, inflating block-validation time and causing full-node synchronisation to lag (Atzei et al., 2017; Chen et al., 2017).[1]

The detection response combined machine-learning-based identification of unusual contract-creation and opcode-invocation patterns with network-level rate limits. The Ethereum Foundation released the Tangerine Whistle hard fork, specified in Ethereum Improvement Proposal 150 (EIP-150) of October 2016, to re-price the affected opcodes, followed by the Spurious Dragon hard fork (EIPs 155, 160, 161, and 170, November 2016) to clear empty account state and address related attack surface. The incident illustrates how protocol-level parameter choices (gas pricing) become security parameters; it also demonstrates the value of empirical monitoring of resource-usage distributions (Kim et al., 2021). Atzei et al. (2017) survey the broader Ethereum smart-contract attack surface of which the 2016 DoS wave is a member, and Chen et al. (Chen et al., 2017) quantify the gas-cost anomalies that made the attacks economically feasible.

Mt. Gox collapse 2014 (exchange-layer). In February 2014, the Tokyo-based cryptocurrency exchange Mt. Gox announced the loss of approximately 850,000 bitcoin (valued at approximately USD 450 million at the time). Transaction malleability, the property that transactions could be re-signed with the same inputs and outputs but a different transaction ID, was implicated in the exchange's accounting failures (Decker and Wattenhofer, 2014; Ron and Shamir, 2013; Böhme et al., 2015). Attackers exploited the fact that the exchange's internal bookkeeping relied on transaction IDs that could be changed after submission, enabling apparent double-withdrawals that were not detected until after the discrepancy had grown beyond recovery.

The detection tools now common for this class of failure, transaction-pattern analysis, cross-system reconciliation, and tighter cryptographic commitment on withdrawal records, were developed and deployed by exchanges in the years following. The incident remains the canonical exchange-layer failure case; its direct causal contribution to Mt. Gox's insolvency is debated, but its role in exposing the transaction-malleability attack surface is well-established (Decker and Wattenhofer, 2014). Böhme et al. (2015) place the collapse within the broader economics and governance literature on Bitcoin and argue that exchange-level institutional design is a first-order determinant of user loss. Bitcoin Improvement Proposal 62 (BIP-62) and subsequent soft-forks mitigated the protocol-layer malleability sources; improved exchange-side accounting closed the application-layer exposure.

Beyond the protocol-layer story, the Mt. Gox collapse was a large-scale consumer-protection failure: many thousands of retail creditors across multiple jurisdictions were unable to access funds for years while civil-rehabilitation proceedings in Japan worked through the estate. That loss of access represented a distinct harm from the theft itself: retail creditors bore litigation costs, currency-value fluctuations across years of insolvency proceedings, and prolonged uncertainty that no on-chain detection system would have prevented. Böhme et al. (2015) place the incident in the wider economics and governance of Bitcoin and argue that exchange-level institutional design, not protocol design, is the first-order determinant of user loss in episodes of this kind. For our framework, the implication is that detection systems focused only on on-chain signals miss an entire class of application-layer risk borne by end users.

[1]Note: we distinguish this attack family from the June 2016 DAO reentrancy exploit, which affected a single smart contract on Ethereum rather than the protocol layer and is documented elsewhere.

Poly Network cross-chain exploit 2021 (cross-chain-layer). In August 2021, approximately USD 610 million in cryptoassets were drained from the Poly Network, a cross-chain interoperability protocol that bridges Ethereum, Binance Smart Chain, and Polygon. The attacker exploited an authorisation logic flaw in the `EthCrossChainManager` contract that allowed the attacker to call the `onlyOwner`-gated `verifyHeaderAndExecuteTx` function with crafted data, promoting their own address to network operator. Most of the funds were subsequently returned after the attacker engaged with the project team, making the economic loss limited in retrospect but the exposure maximal at the time of exploit. Lee et al. (Lee et al., 2023) catalogue the Poly Network compromise together with the 2022 Ronin, Wormhole, and Nomad bridge incidents in their Systematisation-of-Knowledge of cross-chain bridge hacks, and Werner et al. (Werner et al., 2022) situate the general class of bridge failures within a broader DeFi security landscape.

The case illustrates the emerging cross-chain-bridge attack surface and the detection-gap that currently exists for bridges: real-time monitoring of cross-chain message authorisation is less mature than single-chain transaction-anomaly monitoring (Lee et al., 2023; Werner et al., 2022). Subsequent work (Liang et al., 2021) on data-fusion-based collaborative detection in blockchain systems points toward the infrastructure needed to close this gap.

Lessons learned and implications. Three lessons carry across the three cases.

1. **Parameter choices are security choices.** The Ethereum 2016 incident shows that economic parameters (gas prices) carry security consequences; a detection-monitoring regime that did not include opcode-level resource-usage distributions would have been blind to the attack.
2. **Application-layer reconciliation is as important as protocol-layer integrity.** Mt. Gox shows that a protocol-correct ledger can still be paired with an application-layer accounting system whose failure drives the incident.
3. **Cross-chain detection lags single-chain detection.** The Poly Network exploit is representative of a broader class of bridge compromises (Ronin 2022, Wormhole 2022, Nomad 2022) for which the detection literature is thin and the operational tooling is sparse.

These lessons motivate the research agenda in Chapter 6.

6 Future directions for anomaly and fraud detection in blockchain networks

The future of detection work is shaped by the maturation of blockchain as financial infrastructure, the increasing sophistication of adversaries, and the integration of blockchain with traditional systems subject to financial regulation.

Emerging trends and challenges.

1. **Scale and complexity of blockchain networks**. As on-chain transaction counts, contract deployment rates, and cross-chain message volumes grow, detection pipelines must scale accordingly. Streaming architectures and hardware acceleration (Morishima and Matsutani, 2018) are early responses; more are needed.
2. **Adversary sophistication**. Attackers increasingly coordinate across chains, employ mixing services, and exploit DeFi composability to obscure the origin of funds. Detection techniques must evolve correspondingly.

3. **Integration with other systems**. Blockchain interacts with supply-chain, healthcare, and identity systems. Detection must address anomalies that cross system boundaries.
4. **Regulation and compliance**. Growing regulatory attention (the Markets in Crypto-Assets (MiCA) regulation in the EU; the Financial Action Task Force (FATF) travel rule; United States Securities and Exchange Commission (SEC) enforcement) makes detection with audit trails, explainability, and privacy guarantees a regulatory requirement, not an option.

Potential future research directions.

The research frontier for blockchain-fraud detection has in the last three years generated several findings directly relevant to our framework. Qin et al. (2021) provide the empirical baseline for capital-free DeFi attacks, demonstrating that flash-loan composability creates a new exploit class that the detection literature must explicitly address. Qin et al. (2022) quantify the magnitude of maximal-extractable-value (MEV) across Ethereum and show that MEV is a first-order market-integrity concern at the consensus layer, an anomaly class not yet well represented in earlier detection surveys. Zhou et al. (2023) systematise documented DeFi attacks including flash-loan exploits and bridge compromises, providing the classificatory baseline against which our framework can be extended at the contract and cross-chain layers. Sai et al. (2023) apply explainable-AI methods to financial-transaction fraud detection and report that feature-level attribution improves auditability without material accuracy cost, a finding relevant to the XAI research direction listed below.

1. **Machine learning**. Deep learning, reinforcement learning, and graph neural networks remain active research directions. Integration with game-theoretic formalisation (for adversarial robustness) and with XAI methods (for explainability that meets regulatory audit requirements) is a priority.

2. **Network analysis**. Graph-theoretic and social-network-analysis methods applied to transaction and validator graphs continue to produce new detection signals (Li et al., 2019; Maesa et al., 2017).

3. **Game theory**. Mechanism design for protocols that embed detection incentives, and the integration of behavioural-game-theoretic models with statistical detection rules.

4. **Risk assessment**. Formalisation of risk-assessment tooling, continuous monitoring, automated incident response, pre-deployment formal verification, specifically tuned to blockchain applications.

5. **Cross-chain-bridge monitoring**. The Poly Network, Ronin Bridge, Wormhole, and Nomad incidents from 2021-2022 expose the need for real-time authorisation-monitoring tooling for cross-chain messages. Li et al. (2024) provide a 2024 peer-reviewed survey of cross-chain bridge attack surfaces, defences, and open problems that consolidates the literature on this fast-moving subdomain. Even with that systematisation in place, real-time detection at the bridge layer remains a largely open research area.

6. **DeFi flash-loan attack detection**. The detection literature for multi-protocol DeFi exploits (flash-loan-enabled price manipulation, governance-attack constructions) is still emerging.

7. **LLM-based vulnerability detection**. Applying large language models to smart-contract vulnerability detection, at both pre-deployment and runtime, has produced preliminary positive results on public benchmarks; integration with formal verification toolchains remains open.

8. **Privacy-preserving detection**. Detection that operates on encrypted or zero-knowledge-protected data is needed for compliance in jurisdictions with strict data-protection rules.

7 Conclusion

This paper reviewed the detection of anomalies and frauds in blockchain networks, anchored in the financial-market use of the technology. We developed a dimensional classification of anomaly and fraud types (Table 1), a comparative matrix of detection techniques (Table 2), and an analytical framework (Figure 1) that locates each family of techniques in the detection pipeline. We surveyed 31 original peer-reviewed works after removing 4 URL-only entries (plus 24 added for reference currency, case-study verification, and finance-framing) and situated them against three prior general-purpose anomaly-detection surveys.

The central finding is that the detection toolkit is rich at the single-chain single-address level but thin at the cross-chain level and at the interface between on-chain and off-chain systems. Incidents such as the 2021 Poly Network exploit and the 2022 bridge compromises illustrate the practical cost of this gap. The research agenda we outlined in Chapter 6 identifies cross-chain-bridge monitoring, DeFi-flash-loan attack detection, and LLM-based vulnerability detection as priority areas.

Summary of key points.

1. Blockchain networks are financial infrastructure and their security is a financial-stability concern.
2. Anomaly detection in blockchain spans statistical, ML, and game-theoretic methods.
3. Fraud detection additionally requires digital forensics, reputation systems, and risk-assessment frameworks.
4. Case studies, Ethereum 2016 DoS, Mt. Gox 2014, Poly Network 2021, illustrate failure modes at the contract, exchange, and cross-chain layers respectively.
5. The research frontier is in cross-chain monitoring, DeFi exploit detection, XAI for compliance, and privacy-preserving detection.

Implications and recommendations.

For practitioners: adopt a multi-layered detection stack combining statistical baselines, ML-based signature learning, and digital-forensic traceability; maintain logs at the explainability level required by audit and compliance frameworks; stay current with adversary evolution and cross-chain bridge monitoring.

For researchers: close the cross-chain detection gap; develop ML models with XAI-level transparency whose outputs are admissible in forensic and regulatory contexts;

collaborate with industry practitioners and policymakers to ensure research addresses operational constraints.

Limitations. This review has three principal limitations. **First**, this is a narrative review rather than a systematic review with quantitative synthesis, and we therefore make no claims about inter-study effect sizes or pooled estimates. **Second**, the rapidly evolving blockchain ecosystem means that incidents after the inclusion cutoff (2024) are not covered; the research agenda in Chapter 6 is intended to partially offset this by flagging emerging directions. **Third**, we have not independently reproduced the results reported in the cited works; the synthesis relies on the published record.

8 List of Abbreviations

AI artificial intelligence
AML anti-money laundering
API application programming interface
AR autoregressive (model)
ARIMA autoregressive integrated moving average
BAD Blockchain Anomaly Detection
BIP Bitcoin Improvement Proposal
DAO decentralised autonomous organisation
DBSCAN density-based spatial clustering of applications with noise
DeFi decentralised finance
DoS denial of service
EIP Ethereum Improvement Proposal
FATF Financial Action Task Force
GPU graphics processing unit
IQR interquartile range
KYC know your customer
LLM large language model
MA moving average (model)
MEV maximal extractable value
MiCA Markets in Crypto-Assets (EU Regulation)
ML machine learning
PRISMA Preferred Reporting Items for Systematic Reviews and Meta-Analyses
RPC remote procedure call
SARIMA seasonal autoregressive integrated moving average
SEC United States Securities and Exchange Commission
SVM support vector machine
XAI explainable artificial intelligence

Table 1: Dimensional taxonomy of blockchain anomalies and frauds. Each row indexes a representative anomaly or fraud class by the blockchain layer at which it manifests (Network, Consensus, Contract, or Application), the detection method family most applicable to its signature (Statistical, Machine-Learning, Game-Theoretic, or Digital-Forensic), an exemplar real-world incident, and a primary scholarly reference. The taxonomy is intended to serve as a lookup table for practitioners selecting a detection technique for a specific attack class; it is not exhaustive but covers the principal categories reviewed in this paper.

Class	Layer	Detection method	Example incident	Reference
Network outage / partition	Network	Statistical (traffic TS)	Ethereum mempool stalls	Kim et al. 2021
Eclipse / sybil attack	Network	Graph anomaly + ML	Bitcoin eclipse PoC	Akoglu et al. 2015
Selfish mining	Consensus	Game-theoretic	Theoretical, not observed at scale	Baqer et al. 2016
51% attack	Consensus	Statistical + forensic	Ethereum Classic 2020	Ron & Shamir 2013
Reentrancy exploit	Contract	Static + dynamic analysis	DAO 2016	Atzei et al. 2017
Authorization bypass	Contract	Static + ML (code-emb)	Poly Network 2021	Liang et al. 2021
Flash-loan attack	Contract + App	ML (price-TS + graph)	bZx 2020, Harvest 2020	Qin et al. 2021
Rug-pull / exit scam	Contract + App	ML (token-lifecycle)	Numerous DeFi tokens	Mazorra et al. 2022
Double-spending	Consensus + App	Statistical + forensic	Historical Bitcoin attempts	Ron & Shamir 2013
Transaction malleability	Consensus + App	Forensic reconciliation	Mt. Gox 2014	Decker & Wattenhofer 2014
Pump-and-dump	App (market)	ML (price + volume TS)	Various altcoin campaigns	Kamps & Kleinberg 2018
Wash trading	App (market)	Graph + ML	Exchange-level	Mansourifar et al. 2020
Money laundering / mixing	App	Graph + clustering	Multiple	Monamo et al. 2016
Exchange exit scam	App (exchange)	Forensic + off-chain audit	QuadrigaCX 2019	Ron & Shamir 2013
Cross-chain bridge exploit	Cross-chain	Graph + authorization ML	Poly 2021, Ronin 2022	Liang et al. 2021

References

D. Acemoglu, A. Ozdaglar, and A. Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, 2015. doi: 10.1257/aer.20130456.

Table 2: Comparative matrix of anomaly- and fraud-detection techniques for blockchain networks. For each technique family we list the supervision regime (Unsupervised / Semi-supervised / Supervised), the primary data type consumed (Time-series, Graph, Mixed), principal strengths and limitations relevant to financial-infrastructure deployment, and representative scholarly references. This matrix is intended to assist practitioners in matching a detection technique to the data and operational constraints of a specific use case; the strengths and limitations are qualitative summaries of the balance of evidence in the cited literature, not absolute rankings.

Technique	Supervision	Data	Strengths	Limitations	References
ARIMA / SARIMA	Unsupervised	Time-series	Interpretable; low compute	Assumes stationarity; poor on sudden regime shifts	Ahmed et al. 2016
K-means / DBSCAN clustering	Unsupervised	Feature vectors	Label-free; scalable	Sensitive to feature engineering	Pham & Lee 2017
Mahalanobis distance	Unsupervised	Feature vectors	Accounts for covariance	Gaussian assumption; degenerates in high dim	Pham & Lee 2017
Isolation Forest	Unsupervised	Feature vectors	Good high-dim behaviour; fast	Black-box scoring	Shayegan & Sabor 2021
Autoencoder reconstruction	Semi-supervised	Mixed	Learns complex benign manifold	Hyperparameter-sensitive; explainability gap	Taher et al. 2024
Graph Neural Networks (GCN, GAT)	Supervised / Semi	Graph	Captures relational structure	Requires labels or proxies; compute-heavy	Li et al. 2019
Recurrent / transformer sequence models	Supervised	Time-series	Captures long-range dependencies	Label scarcity; overfitting risk	Kim et al. 2021
Ensemble + XAI	Supervised	Mixed	Accuracy + interpretability for compliance	Label requirement; moderate compute	Taher et al. 2024
Bayesian game models	(formal)	N/A	Captures strategic interaction	Hard to parameterise realistically	Baqer et al. 2016
Mechanism design	(formal)	N/A	Preventive; removes incentive	Requires protocol-level change	–
Hash / signature forensic analysis	Supervised	On-chain	Legal-evidence admissibility	Compute for large-window tracing	Ron & Shamir 2013
Transaction-graph forensics	Unsupervised	Graph	Traces fund flows across chains	Breaks on mixers; cross-chain gaps	Maesa et al. 2018; Ober et al. 2013

M. Ahmed, A. N. Mahmood, and R. Islam. A survey of anomaly detection techniques in financial domain. *Future Generation Computer Systems*, 55:278–288, 2015. doi: 10.1016/j.future.2015.01.001.

L. Akoglu, H. Tong, and D. Koutra. Graph based anomaly detection and description: a survey. *Data Mining and Knowledge Discovery*, 29(3):626–688, 2014. doi: 10.1007/s10618-014-0365-y.

N. Atzei, M. Bartoletti, and T. Cimoli. A survey of attacks on Ethereum smart contracts (SoK). In *Principles of Security and Trust (POST 2017)*, volume 10204 of *Lecture Notes in Computer Science*, pages 164–186. Springer, 2017. doi: 10.1007/978-3-662-54455-6_8.

K. Baqer, D. Y. Huang, D. McCoy, and N. Weaver. Stressing out: Bitcoin "stress testing". In *Lecture notes in computer science*, pages 3–18. 2016. doi: 10.1007/978-3-662-53357-4_1.

V. Boginski, S. Butenko, and P. M. Pardalos. Statistical analysis of financial networks. *Computational Statistics & Data Analysis*, 48(2):431–443, 2004. doi: 10.1016/j.csda.2004.02.004.

R. Böhme, N. Christin, B. Edelman, and T. Moore. Bitcoin: Economics, technology, and governance. *Journal of Economic Perspectives*, 29(2):213–238, 2015. doi: 10.1257/jep.29.2.213.

A. G. Bunn, D. L. Urban, and T. H. Keitt. Landscape connectivity: A conservation application of graph theory. *Journal of Environmental Management*, 59(4):265–278, 2000. doi: 10.1006/jema.2000.0373.

Y. Cai and D. Zhu. Fraud detections for online businesses: a perspective from blockchain technology. *Financial Innovation*, 2:20, 2016. doi: 10.1186/s40854-016-0039-4.

V. Chandola, A. Banerjee, and V. Kumar. Anomaly detection: a survey, 2009. URL-only reference, classified as @misc.

T. Chen, X. Li, X. Luo, and X. Zhang. Under-optimized smart contracts devour your money. In *IEEE 24th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2017. doi: 10.1109/saner.2017.7884650.

C. Cholevas, E. Angeli, Z. Sereti, E. Mavrikos, and G. E. Tsekouras. Anomaly detection in blockchain networks using unsupervised learning: A survey. *Algorithms*, 17(5):201–201, 2024. doi: 10.3390/a17050201.

K. Chou. Applications of graph theory to enzyme kinetics and protein folding kinetics. steady and non-steady-state systems, 1990. URL-only reference, classified as @misc.

C. Decker and R. Wattenhofer. Bitcoin transaction malleability and mtgox. In *Computer Security - ESORICS 2014*, volume 8713, pages 313–326. Springer, 2014. doi: 10.1007/978-3-319-11212-1_18.

L. Dobrjanskyj and F. Freudenstein. Some applications of graph theory to the structural analysis of mechanisms. *Journal of Engineering for Industry*, 89(1):153–158, 1967. doi: 10.1115/1.3609988.

P. Gai and S. Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A Mathematical Physical and Engineering Sciences*, 466(2120):2401–2423, 2010. doi: 10.1098/rspa.2009.0410.

Y. Guo and C. Liang. Blockchain application and outlook in the banking industry. *Financial Innovation*, 2:24, 2016. doi: 10.1186/s40854-016-0034-9.

K. M. Han, J. and J. Pei. Data mining: Concepts and techniques, third edition, 2012. URL-only reference, classified as @misc.

N. Hautsch, J. Schaumburg, and M. Schienle. Financial network systemic risk contributions, 2014.

J. Kamps and B. Kleinberg. To the moon: defining and detecting cryptocurrency pump-and-dumps. *Crime Science*, 7(1), 2018. doi: 10.1186/s40163-018-0093-5.

J. Kim, M. Nakashima, W. Fan, S. Wuthier, X. Zhou, I. Kim, et al. Anomaly detection based on traffic monitoring for secure blockchain networking. *2021 IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, pages 1–9, 2021. doi: 10.1109/icbc51069.2021.9461119.

G. Kou, Ö. Olgu Akdeniz, H. Dincer, and S. Yuksel. Fintech investments in european banks: a hybrid IT2 fuzzy multidimensional decision-making approach. *Financial Innovation*, 7:39, 2021. doi: 10.1186/s40854-021-00256-y.

S.-S. Lee, A. Murashkin, M. Derka, and J. Gorzny. SoK: Not quite water under the bridge: Review of cross-chain bridge hacks. In *IEEE International Conference on Blockchain and Cryptocurrency (ICBC)*, 2023. doi: 10.1109/icbc56567.2023.10174993.

N. Li, M. Qi, Z. Xu, X. Zhu, W. Zhou, and S. Wen. Blockchain cross-chain bridge security: Challenges, solutions, and future outlook. *Distributed Ledger Technologies: Research and Practice*, 4(1):1–34, 2024. doi: 10.1145/3696429.

Y. Li, U. Islambekov, C. G. Akçora, E. Smirnova, Y. R. Gel, and M. Kantarcıoğlu. Dissecting ethereum blockchain analytics: What we learn from topology and geometry of ethereum graph, 2019.

W. Liang, L. Xiao, K. Zhang, M. Tang, D. He, and K. Li. Data fusion approach for collaborative anomaly intrusion detection in blockchain-based systems. *IEEE Internet of Things Journal*, 9(16):14741–14751, 2021. doi: 10.1109/jiot.2021.3053842.

Z. Liu, N. C. Luong, W. Wang, D. Niyato, P. Wang, Y.-C. Liang, et al. A survey on applications of game theory in blockchain, 2019.

D. D. F. Maesa, A. Marino, and L. Ricci. Uncovering the bitcoin blockchain: An analysis of the full users graph. pages 537–546, 2016. doi: 10.1109/dsaa.2016.52.

D. D. F. Maesa, A. Marino, and L. Ricci. Data-driven analysis of bitcoin properties: exploiting the users graph. *International Journal of Data Science and Analytics*, 6(1): 63–80, 2017. doi: 10.1007/s41060-017-0074-x.

H. Mansourifar, L. Chen, and W. Shi. Hybrid cryptocurrency pump and dump detection, 2020.

B. Mazorra, V. Adan, and V. Daza. Do not rug on me: Leveraging machine learning techniques for automated scam detection. *Mathematics*, 10(6):949, 2022. doi: 10.3390/math10060949.

P. M. Monamo, V. Marivate, and B. Twala. Unsupervised learning for robust bitcoin fraud detection. pages 129–134, 2016. doi: 10.1109/issa.2016.7802939.

S. Morishima and H. Matsutani. Acceleration of anomaly detection in blockchain using in-gpu cache. pages 244–251, 2018. doi: 10.1109/bdcloud.2018.00047.

E. Nier, J. Yang, T. Yorulmazer, and A. Alentorn. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31(6):2033–2060, 2007. doi: 10.1016/j.jedc.2007.01.014.

M. Ober, S. Katzenbeisser, and K. Hamacher. Structure and anonymity of the bitcoin transaction graph. *Future Internet*, 5(2):237–250, 2013. doi: 10.3390/fi5020237.

N. Paltalidis, D. Gounopoulos, R. Kizys, and Y. Koutelidakis. Transmission channels of systemic risk and contagion in the European financial network. *Journal of Banking and Finance*, 61:S36–S52, 2015. doi: 10.1016/j.jbankfin.2015.03.021.

T. Pham and S. Lee. Anomaly detection in the bitcoin system - a network perspective, 2016a.

T. Pham and S. Lee. Anomaly detection in bitcoin network using unsupervised learning methods, 2016b.

S. Poledna, J. L. Molina-Borboa, S. Martínez-Jaramillo, M. van der Leij, and S. Thurner. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. *Journal of Financial Stability*, 20:70–81, 2015. doi: 10.1016/j.jfs.2015.08.001.

N. Y. Post. Protests in el salvador after bitcoin made official currency. `https://nypost.com/2021/09/16/protests-in-el-salvador-after-bitcoin-made-official-currency/`, 2021. URL-only reference, classified as @misc.

K. Qin, L. Zhou, B. Livshits, and A. Gervais. Attacking the DeFi ecosystem with flash loans for fun and profit. In *Financial Cryptography and Data Security (FC)*, 2021. doi: 10.1007/978-3-662-64322-8_1.

K. Qin, L. Zhou, and A. Gervais. Quantifying blockchain extractable value: How dark is the forest? In *IEEE Symposium on Security and Privacy (S&P)*, 2022. doi: 10.1109/sp46214.2022.9833734.

D. Ron and A. Shamir. Quantitative analysis of the full bitcoin transaction graph. In *Lecture notes in computer science*, pages 6–24. 2013. doi: 10.1007/978-3-642-39884-1_2.

C. V. Sai, D. Das, N. S. Elmitwally, O. Elezaj, and M. B. Islam. Explainable AI-driven financial transaction fraud detection using machine learning and deep neural networks. SSRN preprint, 2023.

S. Sayadi, S. B. Rejeb, and Z. Choukair. Anomaly detection model over blockchain electronic transactions, 2019.

M. J. Shayegan and H. R. Sabor. A collective anomaly detection method over bitcoin network, 2021.

M. Signorini, M. Pontecorvi, W. Kanoun, and R. D. Pietro. Bad: A blockchain anomaly detection solution. *IEEE Access*, 8:173481–173490, 2020. doi: 10.1109/access.2020.3025622.

S. S. Taher, S. Y. Ameen, and J. A. Ahmed. Advanced fraud detection in blockchain transactions: An ensemble learning and explainable ai approach. *Engineering Technology & Applied Science Research*, 14(1):12822–12830, 2024. doi: 10.48084/etasr.6641.

P. Vanini, S. Rossi, E. Zvizdic, and T. Domenig. Online payment fraud: from anomaly detection to risk management. *Financial Innovation*, 9:66, 2023. doi: 10.1186/s40854-023-00470-w.

S. M. Werner, D. Pérez, L. Gudgeon, A. Klages-Mundt, D. Harz, and W. J. Knottenbelt. SoK: Decentralized finance (DeFi). In *4th ACM Conference on Advances in Financial Technologies (AFT '22)*, 2022. doi: 10.1145/3558535.3559780.

P. Xia, H. Wang, B. Zhang, R. Ji, B. Gao, L. Wu, X. Luo, and G. Xu. Characterizing cryptocurrency exchange scams. *Computers & Security*, 98:101993, 2021. doi: 10.1016/j.cose.2020.101993.

M. Xu, X. Chen, and G. Kou. A systematic review of blockchain. *Financial Innovation*, 5:27, 2019. doi: 10.1186/s40854-019-0147-z.

L. Zhang, X. Ma, and Y. Liu. Sok: Blockchain decentralization. *arXiv preprint arXiv:2205.04256*, 2023. arXiv:2205.04256.

R. Zhang, G. Zhang, L. Liu, C. Wang, and S. Wan. Anomaly detection in bitcoin information networks with multi-constrained meta path. *Journal of Systems Architecture*, 110:101829–101829, 2020. doi: 10.1016/j.sysarc.2020.101829.

L. Zhou, X. Xiong, J. Ernstberger, S. Chaliasos, Z. Wang, Y. Wang, K. Qin, R. Wattenhofer, D. Song, and A. Gervais. SoK: Decentralized finance (DeFi) attacks. In *IEEE Symposium on Security and Privacy (S&P)*, 2023. doi: 10.1109/sp46215.2023.10179435.